



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 540 161 A1

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 92308365.3

(51) Int. Cl.⁵: G06F 9/46

(22) Date of filing: 15.09.92

(30) Priority: 01.11.91 GB 9123264

(43) Date of publication of application:
05.05.93 Bulletin 93/18

(84) Designated Contracting States:
BE DE FR GB IT

(71) Applicant: INTERNATIONAL COMPUTERS
LIMITED
ICL House
Putney, London, SW15 1SW(GB)

(72) Inventor: Fields, Michael
37 Victoria Street
Ramsbottom, Lancashire(GB)
Inventor: Holt, Nicholas Peter
104 Padfield Main Road, Padfield
Hadfield via Hyde, Cheshire SK14 7ET(GB)
Inventor: Fullerton, Mark Norman
7 Simmondley Grove
Glossop, Derbyshire(GB)
Inventor: Knowles, Andrew John
6 East Vale
Marple, Cheshire(GB)

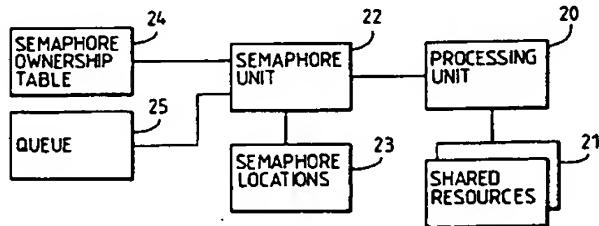
(74) Representative: Guyatt, Derek Charles Patents
and Licensing International Computers
Limited et al
Six Hills House London Road
Stevenage, Herts, SG1 1YB (GB)

(5) Semaphore mechanism for data processing system.

(57) In a data processing system a number of processing nodes share resources. Access to the shared resources is controlled by semaphores, each node having a local copy of all the semaphores. Nodes may acquire ownership of semaphores. When a node requires a semaphore operation on a particular semaphore, a semaphore message is broadcast to all the nodes instructing them to perform the semaphore operation on their local copies of the semaphore. If the semaphore is unowned, the node must suspend the semaphore operation until the message returns, so as to ensure correct chronology for the semaphore operation. If, however, the semaphore owned by this node, the node can perform the semaphore operation without waiting for the message to return. This speeds up the semaphore mechanism. If the semaphore is owned by another

node, that other node relinquishes ownership so that the semaphore operation can be performed.

Fig. 2.



EP 0 540 161 A1

Background to the Invention

This invention relates to semaphore mechanisms for data processing systems.

More specifically, the invention is concerned with a data processing system having a number of processing nodes, interconnected by a communication network, wherein the nodes have access to shared resources such as shared memory.

A semaphore mechanism is a means for ensuring that only one processing node can access a shared resource at any given time. In one known form of semaphore mechanism, each shared resource (eg each area of the shared memory) has a semaphore location associated with it. Whenever a node wishes to access a shared resource, it first increments and tests the associated semaphore location, to find out whether another node is already accessing that resource. Conversely, when the node has completed its access to a shared resource, it tests and decrements the appropriate semaphore location. Each node holds local copies of all the semaphore locations, and to ensure that all these copies are kept consistent, each semaphore action requires sending a semaphore message over the network to initiate the required operation on each local copy of the semaphore location.

In such a system, it is necessary to provide some way of synchronising semaphore operations, to ensure that each node applies all such operations in the same chronological sequence, regardless of which node originated them. One way of achieving this is to broadcast the semaphore messages to all the nodes in the system, including the originating nodes. These messages are distributed to the nodes in such a way that all the nodes receive the messages in the same sequence. This sequence then defines a unique chronology for the semaphore operations.

A problem with this is that a process running in a node must be suspended each time it initiates a semaphore operation, until the associated semaphore message has been "turned round" (i.e. received back from the network) so that the process can be sure of the correct chronology for the operation. This can result in significant delays in performing the semaphore operations.

The object of the present invention is to provide a way of reducing the suspension time associated with semaphore operations, and hence to speed up the operation of the system.

Summary of the Invention

According to the invention there is provided a data processing system comprising :

(a) a plurality of processing nodes,

- (b) a communications network interconnecting the nodes to allow the nodes to send messages over the network, the messages being delivered to all the nodes in the same chronological sequence,
- (c) a plurality of shared resources, shared by all the nodes,
- (d) a plurality of semaphores for controlling access to respective shared resources, each node having a local copy of all the semaphores,
- (e) means for defining ownership states for the semaphores whereby each semaphore may be unowned, or owned by a particular node,
- (f) means in each node, responsive to a request for a semaphore operation on an unowned semaphore, for sending a semaphore message over the network instructing all the nodes to perform the semaphore operation on their local copies of the semaphore, the semaphore operation being suspended in this node until the semaphore message has been returned from the network, and
- (g) means in each node, responsive to a request for a semaphore operation on a semaphore owned by this node, for sending a semaphore message over the network instructing all the nodes to perform the semaphore operation on their local copies of the semaphore, this node being permitted to proceed with the semaphore operation without waiting for the semaphore message to be returned from the network.

It will be seen that, as a result, the number of times a node has to suspend while waiting for a semaphore message to be synchronised is greatly reduced. An additional delay will occur when two nodes attempt to use a particular semaphore at the same time, but in practice this situation arises relatively infrequently and hence this is a relatively small overhead problem.

Brief Description of the Drawings

Figure 1 is a block diagram of a data processing system embodying the invention.

Figure 2 is a block diagram showing a processing node of the system in more detail.

Figures 3 - 12 are flow diagrams illustrating the semaphore mechanism with each node.

Description of an Embodiment of the Invention

One data processing system in accordance with the invention will now be described by way of example with reference to the accompanying drawings.

Referring to Figure 1, the data processing system comprises a plurality of processing nodes 10, interconnected by a communications network

12. Each node can send messages over the network to the other nodes. The network broadcasts each message to all the nodes, including the node that originated the message. The network is a star coupler, which ensures that all the nodes receive the messages in the same sequence. This sequence defines a unique chronology for the messages.

Referring to Figure 2, each node includes a processing unit 20 for scheduling and executing a number of concurrent processes. The processing unit may be conventional and so will not be described in further detail.

The processing unit 20 in each node has access to a number of shared resources 21, ie resources that are shared with other nodes. These resources may, for example, comprise areas of shared memory.

Each node includes a semaphore unit 22 which controls access to the shared resources 21, using a set of semaphore locations 23, a semaphore ownership table 24 and a semaphore queue 25. Each node has its own local copies of the semaphore locations and the semaphore ownership table, and has its own semaphore queue.

Each shared resource has a particular semaphore location associated with it. If the resource is free, the semaphore location has the value -1.

Whenever a process wishes to gain access to a shared resource, it issues an INCT (increment and test) request. The semaphore location associated with that resource is then incremented, and its value returned. If the returned semaphore value is 0, the operation is successful, and the process may access the shared resource. If the returned semaphore value is greater than 0, the operation is unsuccessful.

When the process has finished using the shared resource, it issues a TDEC (test and decrement) request. This causes the value of the semaphore location to be returned, and the location is then decremented. If the returned value is 0, this means that the shared resource is now free, and is available for access by another process.

Both of these operations (INCT and TDEC) are atomic, in the sense that they may not be interrupted by another operation.

The semaphore ownership table 24 consists of a number of sections, one for each node. Each section has a fixed number of slots, each of which can hold an entry, defining the ownership state of a particular semaphore. Each entry in the ownership table comprises the following fields:

ADDRESS: the virtual address of the semaphore location to which the entry relates.

STATE: the ownership state of the semaphore location (see below).

Each semaphore location has one of the following ownership states: IDLE, OWNED, DISCARD, OTHER-OWNED, and QUEUED. The current state of a semaphore is indicated by the STATE field in the corresponding entry in the ownership table. The exception to this is the IDLE state, which represents the state where there is no entry for the semaphore in the ownership table.

The meanings of these ownership states are as follows:

IDLE: the semaphore is not owned by any node. OWNED: the semaphore is owned by the local node. DISCARD: ownership of the semaphore has been relinquished. OTHER-OWNED: the semaphore is owned by a remote node. QUEUED: the semaphore has one or more suspended semaphore operations in the queue 25.

The semaphore queue 25 is a list of suspended semaphore operations. As will be described, whenever a node requires to perform a semaphore operation on a semaphore currently owned by another node, that operation is placed in the queue. The operation is applied when the owning node relinquishes ownership.

Each entry in the queue comprises the following fields:

FUNCTION: the required semaphore operation (INCT or TDEC). ADDRESS: the virtual address of the semaphore location. VALUE: the expected value of the semaphore location (ie the value of the semaphore when the semaphore message was turned round).

The semaphore units 22 in the different nodes communicate with each other by means of semaphore messages sent over the network. Each message comprises the following fields:

FUNCTION: the type of message (see below).

ORIGIN: the identity of the source node.

ADDRESS: the virtual address of the semaphore location to which the message relates.

READ-VALUE: the value of the semaphore location read by the originating node before sending the message.

The FUNCTION field represents one of the following functions:

INCT: increment and test specified semaphore.

TDEC: test and decrement a specified semaphore.

APPLIED-INCT: as for INCT, except that the INCT has already been applied by the originating node.

APPLIED-TDEC: as for TDEC, except that the TDEC has already been applied by the originating node.

DISCARD: used by the owner of a semaphore to relinquish ownership.

It should be noted that the READ-VALUE field is not required for DISCARD messages.

General Rules

The general rules of operation of the semaphore mechanism are as follows.

(1) Message ordering

All the messages pertaining to semaphore operations (Semops) are globally ordered. Thus a message will be received in all nodes at the same point in the sequence.

(2) Assumption of Ownership

A node may only assume ownership after issuing an INCT operation. It must then suspend the issuing process. Once the INCT operation has been turned round by the network, the node examines the ownership table. If no other node has earlier assumed ownership then the originating node may do so. Other nodes upon encountering the INCT will have received the same sequence of earlier messages and are thus capable of reproducing the originating node's decision and updating their local copies of the table accordingly. A node thus assumes ownership at a particular point in the sequence of messages.

(3) Relinquishing Ownership

A node may relinquish ownership of a semaphore owned by it, by sending either TDEC or a DISCARD message.

A node relinquishes ownership at a particular point in the sequence. The node which issues the relinquish regards it as no longer owned as soon as it issues the TDEC or DISCARD message. Ownership is not, however, fully discarded until the message has been turned round by the network.

(4) Discarding when Ownership Table Full

The ownership table in each node is obviously finite. The case therefore occurs that the table is full but the node wishes to issue another INCT operation. In this case the node must choose one of its owned semaphores and relinquish it (ie issue a DISCARD operation on it), thus freeing a slot of the ownership table. In order that there will always be at least one slot which can be discarded, the table must be larger than the total number of semaphores which can be outstanding at one time, for example, if the system can have only one semaphore per node outstanding at once, a minimum of two slots are required for each node.

(5) Alien Semops

When a node receives a semop from an alien node to a location which it owns, it must issue a DISCARD message to allow the alien node to complete its operation. It must not ignore the request in the expectation of later sending a TDEC message out. It is however, permitted to wait for a short time before issuing a DISCARD, in the hope that a TDEC will be generated.

(6) Gain Ownership on INCT from -1 to 0 only

In order to improve the match between the control of a shared resource and the ownership of the semaphore that manages the resource, it is desirable to allow a node to gain ownership of a semaphore only when control of the resource is gained. As the gaining of control is under software control it is impossible to produce a perfect match. However, if control is gained only when a semaphore location is incremented from -1 to 0, this gives a good approximation to the ideal behaviour.

Outgoing INCT

Referring to Figure 3, this shows the action of the semaphore unit 22 when it receives an INCT request from a process running on the local processing unit.

The INCT request includes as parameters the virtual address of the semaphore location to which it relates and also the current value of that location. The process that issued the request is suspended, waiting for the semaphore operation to be synchronised.

The semaphore unit first searches (31) its own section of the ownership table 24 (ie the section relating to the local node) looking for an entry in the OWNED state.

If such an entry is found, this means that the semaphore is already owned by this node, and so the node can proceed to perform the semaphore operation without waiting for the operation to be synchronised. An APPLIED-INCT message is sent (32) over the network, this message including the semaphore address and value as supplied with the INCT request. The INCT operation is applied (33) to the local copy of the semaphore and the returned value is passed (34) to the processing unit. The process that initiated the INCT can then be de-suspended.

If, on the other hand, the semaphore is not owned by this node, an INCT message is sent (35) over the network, this message containing the semaphore address and value as supplied by the INCT request. The process that initiated the INCT request remains suspended, waiting for the mes-

sage to be turned around.

As will be described, when a node sends an INCT message, it may acquire ownership of the semaphore, if that semaphore is not already owned by another node. Therefore, before sending the INCT message, the semaphore unit first checks (36) the semaphore ownership table to determine whether there are any free entries in its own section of the table. If this section of the table is full, one of the existing entries is selected and a DISCARD message is sent (37) over the network informing all the nodes that this entry is to be discarded. The state of the selected entry is set to DISCARD.

Outgoing TDEC

Referring to Figure 4, this shows the action of the semaphore unit when it receives a TDEC request from a process running on the local processing unit.

The TDEC request includes as parameters the virtual address of the semaphore location to which it relates and also the current value of that location. The process that issued the request is suspended, waiting for the semaphore operation to be synchronised.

The semaphore unit first searches (41) its own section of the ownership table, looking for an entry in the OWNED state.

If such an entry is found, this means that the semaphore is already owned by this node, and so the node can proceed to perform the semaphore operation without waiting for the operation to be synchronised. An APPLIED TDEC message is sent (42) over the network, this message including the semaphore address and value as supplied with the TDEC request. The TDEC operation is applied (43) to the local copy of the semaphore, and the returned value is passed (44) to the processing unit. The process that initiated the TDEC can then be de-suspended. The table entry is set (45) to DISCARD to indicate that this node has now relinquished ownership of the semaphore.

If, on the other hand, the semaphore is not owned by this node, a TDEC message is sent (46) over the network, this message containing the semaphore address and value, as supplied by the TDEC request.

Incoming own INCT

Referring to Figure 5, this shows the operation of the semaphore unit when the node receives an INCT message from the network, originating from itself.

The semaphore unit first searches (51) the whole of its ownership table, looking for an entry

relating to the semaphore specified in the INCT message.

If no entry can be found (ie the semaphore is IDLE), the semaphore value in the INCT message is tested (52) and, if it is equal to -1, the node assumes ownership of the semaphore, by creating (53) an entry in its own section of the ownership table, setting the state of this entry to OWNED. The INCT operation is then applied (54) to the local copy of the semaphore, and the returned value is passed (55) to the processing unit. The process that initiated the INCT can then be de-suspended.

If an OWNED or DISCARD entry is found (56), the INCT is applied (57) to the local copy of the semaphore and the returned value is passed (58) to the processing unit, so that the process can be de-suspended.

If an OTHER-OWNED or QUEUED entry is found (59), the INCT is applied (510) to the local copy of the semaphore, but the returned value is not passed to the processor. Instead, an INCT entry is added (511) to the local queue, this entry containing the semaphore address and returned value. If the semaphore is OTHER-OWNED, its state is changed to QUEUED.

Incoming own TDEC

Referring to Figure 6, this shows the operation of the semaphore unit when the node receives a TDEC message from the network, originating from itself.

The semaphore unit first searches (61) the whole of its ownership table, looking for an entry relating to the semaphore specified in the TDEC message.

If no entry can be found (ie the semaphore is IDLE), the TDEC operation is applied (62) to the local copy of the semaphore, and the returned value is passed (63) to the processing unit. The process that initiated the TDEC can then be de-suspended.

If an OWNED or DISCARD entry is found (64), the TDEC is applied (65) to the local copy of the semaphore and the returned value is passed (66) to the processing unit, so that the process can be de-suspended.

If an OTHER-OWNED or QUEUED entry is found (67), the TDEC is applied (68) to the local copy of the semaphore, and a TDEC entry is added (69) to the local queue, this entry containing the semaphore address and returned value. If the semaphore is OTHER-OWNED, its state is changed to QUEUED.

Incoming own APPLIED - INCT

When a node receives an incoming APPLIED - INCT message originating from itself, the semaphore unit acts as follows.

The semaphore unit first searches its own section of the ownership table, looking for an entry relating to the specified semaphore. If an entry in the DISCARD or OWNED state is found, no action is required. Otherwise, an error is indicated.

Incoming own APPLIED - TDEC or DISCARD

Referring to Figure 7, this shows the action of the semaphore unit when the node receives an APPLIED - TDEC or DISCARD message that originated from itself.

The semaphore unit first searches (71) its own section of the ownership table, looking for an entry for the specified semaphore. If an entry in the DISCARD state is found (72), that entry is deleted (73) so as to reset the state of the semaphore to IDLE). Otherwise, an error is indicated.

Incoming Alien INCT

Referring to Figure 8, this shows the action of the semaphore unit when it receives an INCT message originating from another node.

The semaphore unit first searches (81) the whole of the ownership table looking for an entry relating to the specified semaphore.

If no entry is found (ie the semaphore is IDLE), the INCT operation is applied (82) to the local copy of the semaphore. Also the semaphore value contained in the message is tested (83) and, if it is -1, an entry is created (84) in the section of the ownership table corresponding to the node that originated the message, with state OTHER - OWNED to indicate that this semaphore is now owned by that node.

If an OWNED entry is found (85), the INCT operation is applied (86) to the local copy of the semaphore. A DISCARD message is then sent (87) and the state of the entry is set (88) to DISCARD.

If a DISCARD, OTHER - OWNED or QUEUED entry is found (89), the INCT operation is applied (810) to the local copy of the semaphore.

Incoming Alien TDEC

Referring to Figure 9, this shows the action of the semaphore unit when it receives a TDEC message originating from another node.

The semaphore unit first searches (91) the whole of the ownership table looking for an entry relating to the specified semaphore.

If no entry is found (ie the semaphore is IDLE), then the TDEC operation is applied (92) to the local copy of the semaphore.

- 5 If an OWNED entry is found (93), the TDEC operation is applied (94) to the local copy of the semaphore. A DISCARD message is then sent (95), and the state of the entry is set (96) to DISCARD.
- 10 If a DISCARD, OTHER - OWNED or QUEUED entry is found (97), the TDEC operation is applied (98) to the local copy of the semaphore.

Incoming Alien APPLIED - INCT

- 15 Referring to Figure 10, this shows the action of the semaphore unit when it receives an APPLIED - INCT message originating from another node.

The semaphore unit first searches (101) the section of the ownership table relating to the node that originated the message, looking for an entry relating to the specified semaphore. If no entry is found then an error is indicated.

- 20 If an OTHER - OWNED entry is found (102), then the INCT operation is applied (103) to local copy of the semaphore.

If a QUEUED entry is found (104), then the INCT operation is applied (105) to the local copy of the semaphore. Each entry in the local queue associated with this particular semaphore is then adjusted (106) by adding 1 to its value field.

If the semaphore is in any other state, an error is indicated.

Incoming Alien APPLIED - TDEC

- 35 Referring to Figure 11, this shows the action of the semaphore unit when it receives an APPLIED - TDEC message originating from another node.

The semaphore unit first searches (111) the section of the ownership table relating to the node that originated the message, looking for an entry relating to the specified semaphore. If no entry is found, an error is indicated.

- 40 If an OTHER - OWNED entry is found (112), the TDEC operation is applied (113) to the local copy of the semaphore, and the entry is deleted (114) so as to set the semaphore to IDLE.

If a QUEUED entry is found (115), the TDEC operation is applied (116) to the local copy of the semaphore. Each entry in the local queue associated with this semaphore is then adjusted (117) by subtracting 1 from its value field. The queue is then applied (118) by passing each entry associated with this semaphore to the processing unit. All these queue entries are then deleted from the queue, and the remaining entries are shifted down to fill up the gaps. Finally, the entry in the ownership table is deleted (119) so as to set the sema-

phore to IDLE.

Incoming Alien DISCARD

Referring to Figure 12 this shows the action of the semaphore unit when it receives a DISCARD message originating from another node.

The semaphore unit first searches (121) the section of the ownership table relating to the node that originated the message, looking for an entry relating to the specified semaphore. If no entry is found, an error is indicated.

If an OTHER-OWNED entry is found (122), this entry is deleted (123) so as to set the semaphore to IDLE.

If a QUEUED entry is found (124), the local queue is applied (125) by passing each entry in this queue associated with this semaphore to the processing unit. All these entries are deleted from the queue and the remaining entries are shifted down to fill up the gaps. Finally, the entry in the ownership table is deleted (126) so as to set the semaphore to IDLE.

Example of Operation (1)

The case will first be considered where a process running in node 0 requires to access a shared resource, where the semaphore for that resource is not currently owned by any node (ie the semaphore is IDLE).

The process issues an INCT request to the semaphore unit 22 and is then suspended. Referring to Figure 3, the semaphore unit does not find an entry for this semaphore, and so sends an INCT message over the network to all the nodes.

As shown in Figure 8, when this INCT message is received at each other node the INCT is applied to the local copy of the semaphore, and an OTHER-OWNED entry is created in the ownership table in that node, indicating that the semaphore is now owned by node 0.

As shown in Figure 5, when the INCT message is received at node 0, an OWNED entry is created in the ownership table, indicating that the semaphore is now owned by this node. The INCT is then applied to the local copy of the semaphore and the returned value is passed to the processor. The suspended process can then be de-suspended and can access the shared resource.

When the process has completed its access to the shared resource, it issues a TDEC request and is suspended again. Referring to Figure 4, the semaphore unit finds that the semaphore is already owned by this node, and so can immediately apply the TDEC operation to the local copy of the semaphore, passing the returned value to the processor. The suspended process can therefore be

immediately de-suspended, without waiting for the TDEC message to be turned around. Thus, the semaphore operation is effectively applied ahead of its true chronological order. The semaphore unit also sends an APPLIED-TDEC message over the network and sets the state of the entry in the ownership table to DISCARD.

Referring to Figure 11, when the APPLIED-TDEC message is received in the other nodes, the semaphore unit in each of these nodes finds that the semaphore is in the OTHER-OWNED state. It therefore applies the TDEC to its local copy of the semaphore and deletes the entry in the ownership table (ie sets the semaphore to IDLE).

Referring to Figure 7, when the APPLIED-TDEC message is received by node 0, the semaphore unit in that node finds that the semaphore is in the DISCARD state. It therefore deletes the entry in the ownership table, returning the semaphore to the IDLE state.

Example of the Operation (2)

Consider now the case where a process running in node 0 requires to access a shared resource, where the semaphore for that resource is already owned by node 1.

As before, the process issues an INCT request, and is suspended, and the semaphore unit sends an INCT message (Figure 3).

Referring to Figure 8, when the INCT message is received at node 1, the semaphore unit in node 1 finds that the semaphore is OWNED, ie owned by node 1. It therefore applies the INCT to the local copy of the semaphore, and then sends a DISCARD message, setting the ownership state of the semaphore to DISCARD. Thus, node 1 relinquishes ownership of the semaphore to allow node 0 to perform its semaphore operation.

Referring to Figure 5, when the INCT message is received at node 0, the semaphore unit in node 0 finds that the semaphore is OTHER-OWNED. It therefore applies the INCT to the local copy of the semaphore, and puts the returned value into the local queue, setting the ownership state of the semaphore to QUEUED.

Referring to Figure 12, when the DISCARD message from node 1 arrives at node 0, the semaphore unit finds that the semaphore is QUEUED. It therefore applies the queue, passing the queued value to the processor so that the process can now be de-suspended. Finally, it deletes the local ownership table entry, returning the semaphore to the IDLE state.

When the DISCARD message arrives at node 1, the local ownership table entry is deleted.

Claims

1. A data processing system comprising:

- (a) a plurality of processing nodes,
- (b) a communications network interconnecting the nodes to allow the nodes to send messages over the network, the messages being delivered to all the nodes in the same chronological sequence,
- (c) a plurality of shared resources, shared by all the nodes, and
- (d) a plurality of semaphores for controlling access to respective shared resources, each node having a local copy of all the semaphores,

characterised by:

- (e) means for defining ownership states for the semaphores whereby each semaphore may be unowned, or owned by a particular node,
- (f) means in each node, responsive to a request for a semaphore operation on an unowned semaphore, for sending a semaphore message over the network instructing all the nodes to perform the semaphore operation on their local copies of the semaphore, the semaphore operation being suspended in this node until the semaphore message has been returned from the network, and
- (g) means in each node, responsive to a request for a semaphore operation on a semaphore owned by this node, for sending a semaphore message over the network instructing all the nodes to perform the semaphore operation on their local copies of the semaphore, this node being permitted to proceed with the semaphore operation without waiting for the semaphore message to be returned from the network.

2. A system according to Claim 1 wherein each node includes means, responsive to a request for a semaphore operation on a semaphore owned by another node, for sending a semaphore message over the network instructing all the nodes to perform the semaphore operation on their local copies of the semaphore, the semaphore operation being suspended until permission to proceed is received from that other node.

3. A system according to Claim 2 wherein said other node gives permission to proceed by sending a message over the network indicating that it is relinquishing ownership of the semaphore.

4. A system according to Claim 2 or 3 wherein each node includes a queue for queuing values returned by semaphore operations on semaphores owned by other nodes.

5. A system according to any preceding claim wherein each node includes means for acquiring ownership of an unowned semaphore when a semaphore message sent by this node has been returned from the network.

6. A system according to any preceding claim wherein each node relinquishes ownership of a semaphore owned by it when it performs a semaphore operation to free the associated shared resource, or when it receives a semaphore message relating to that semaphore from another node.

7. A system according to any preceding claim wherein the means for defining ownership states for semaphores comprises an ownership table, each node having a local copy of the ownership table.

Fig. 1.

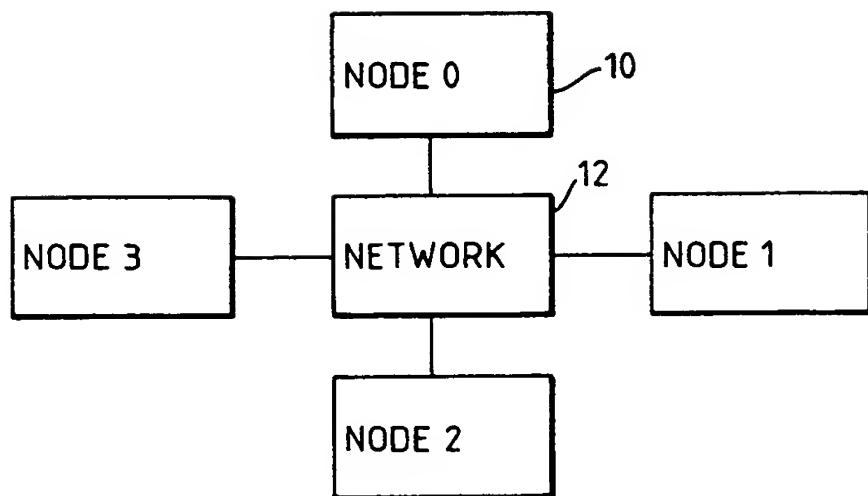


Fig. 2.

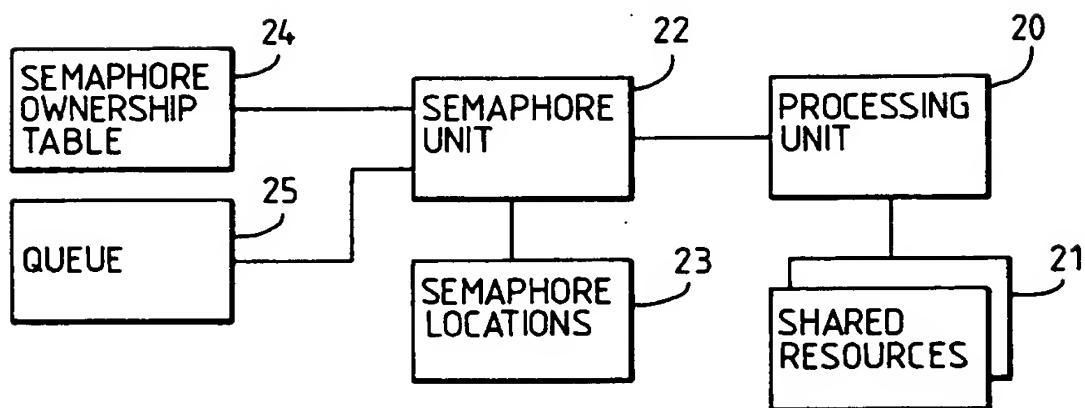


Fig. 3.

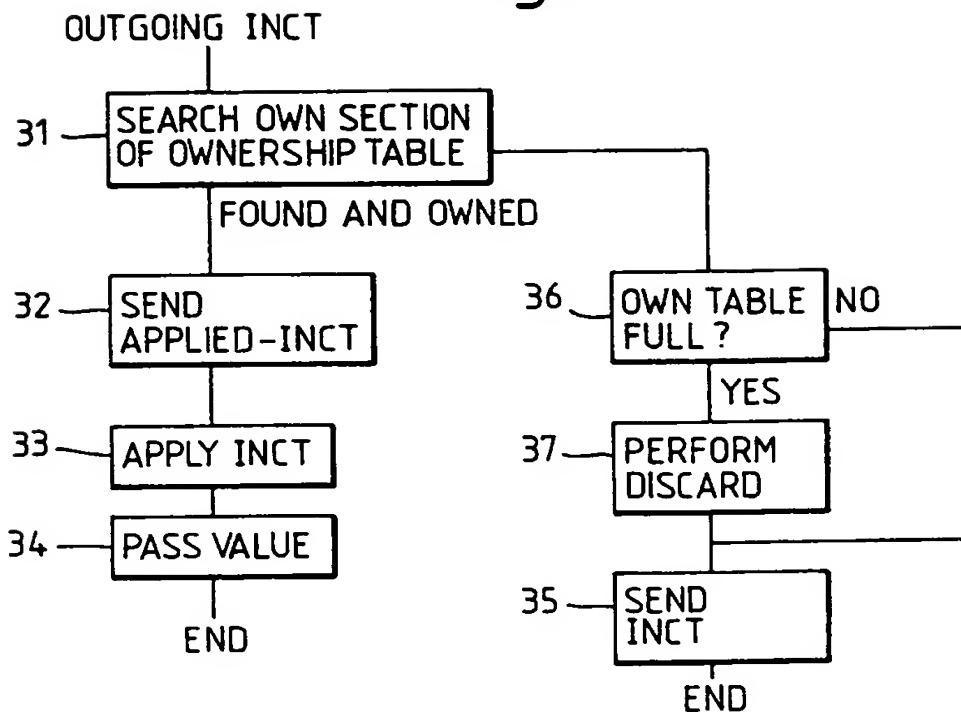


Fig. 4.

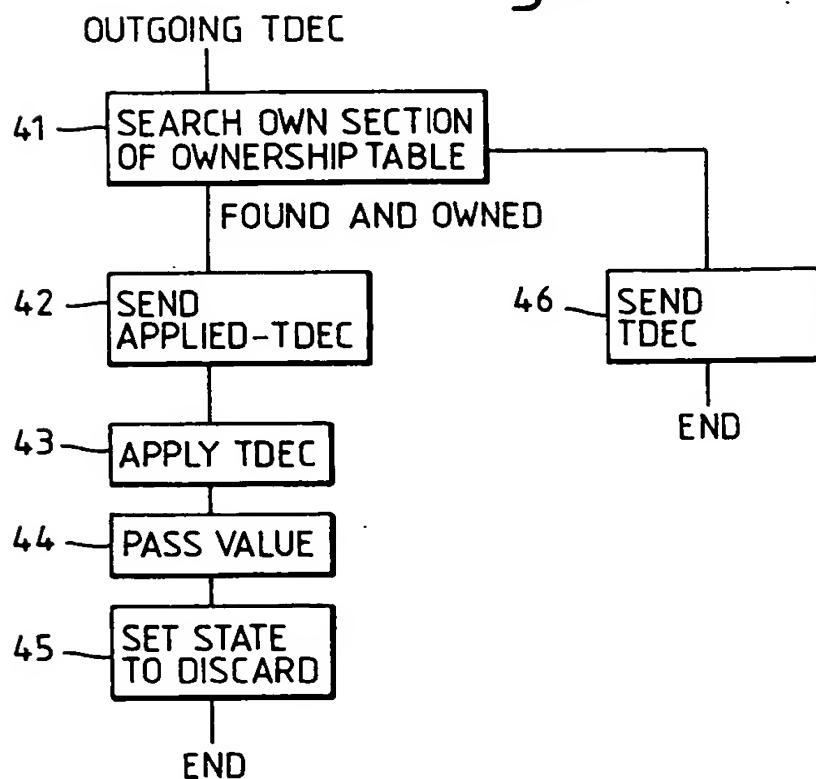


Fig. 5.

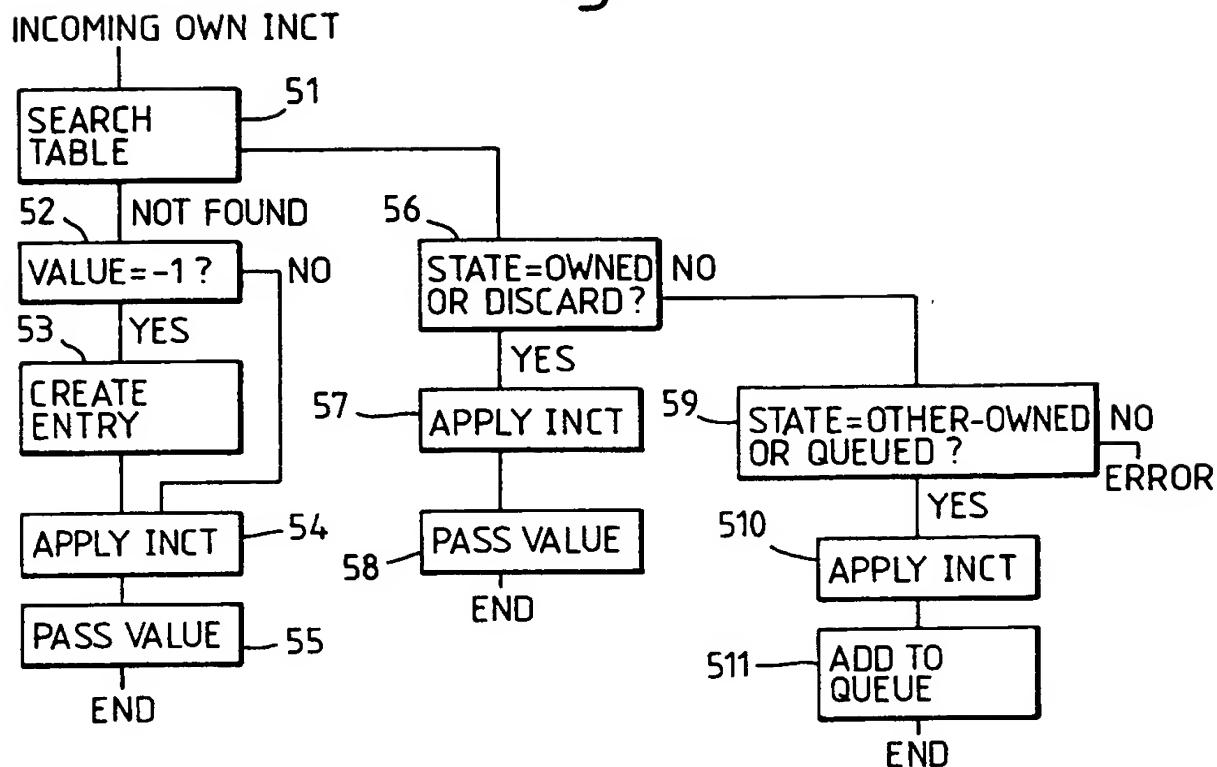


Fig. 6.

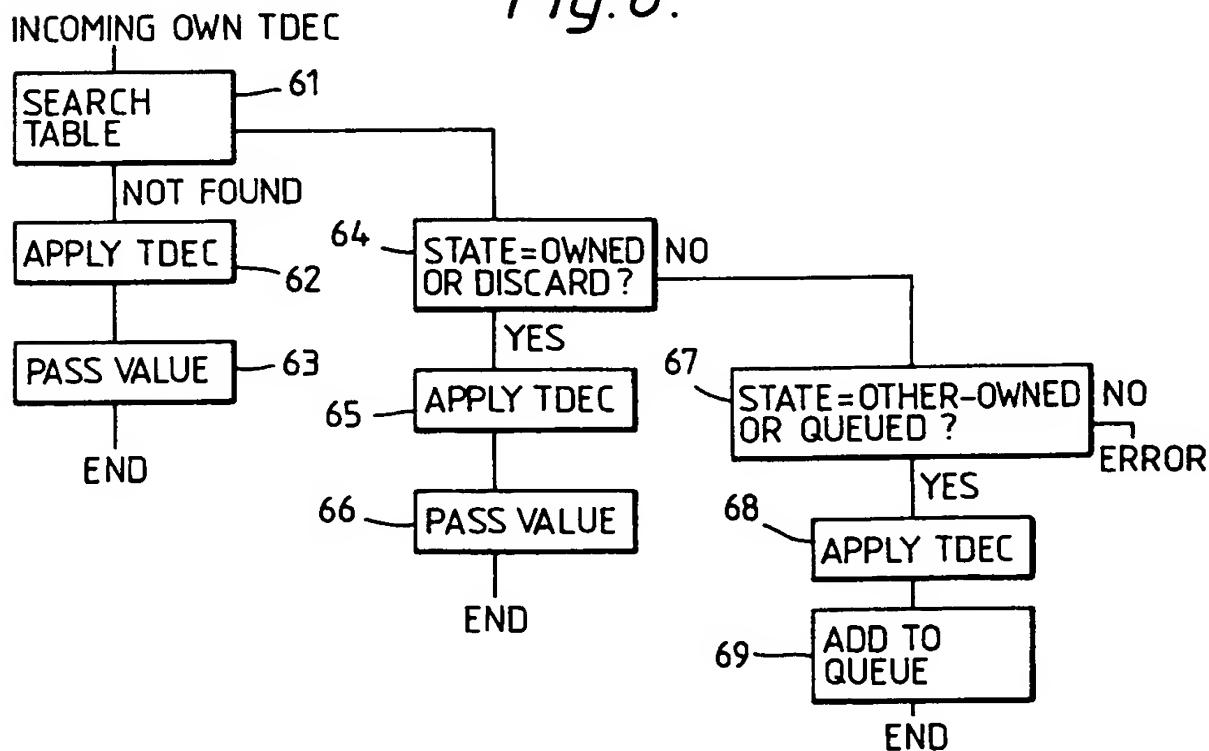


Fig. 7.

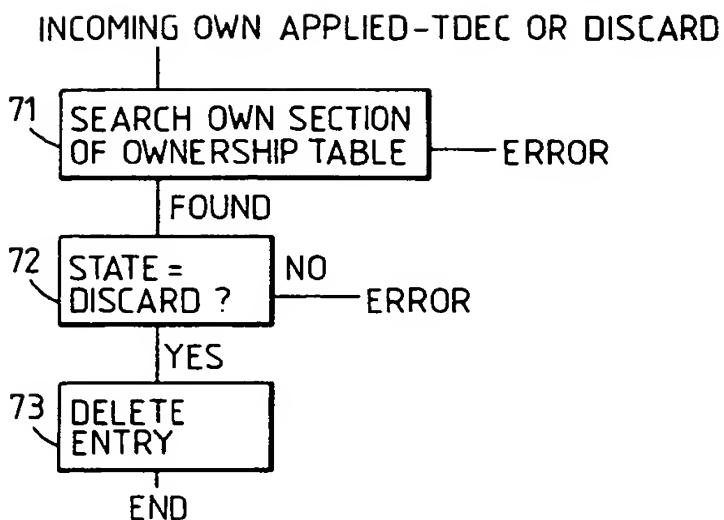
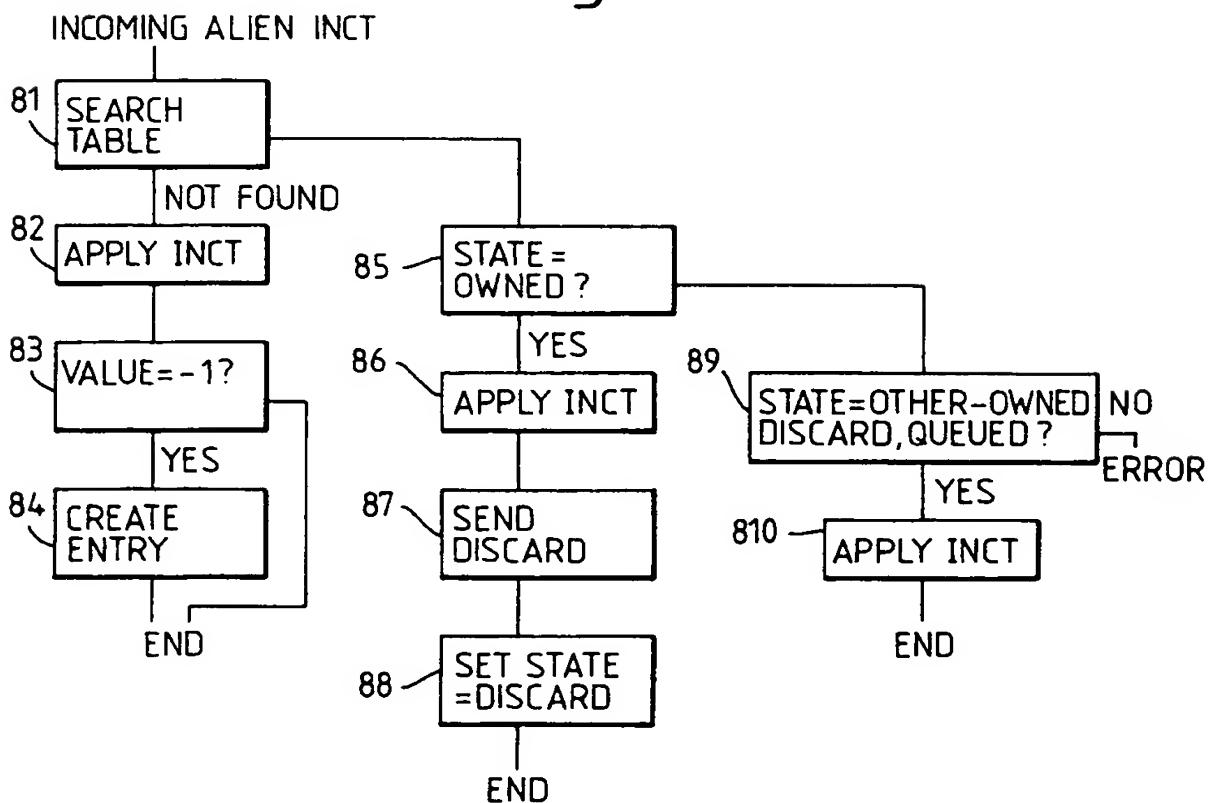
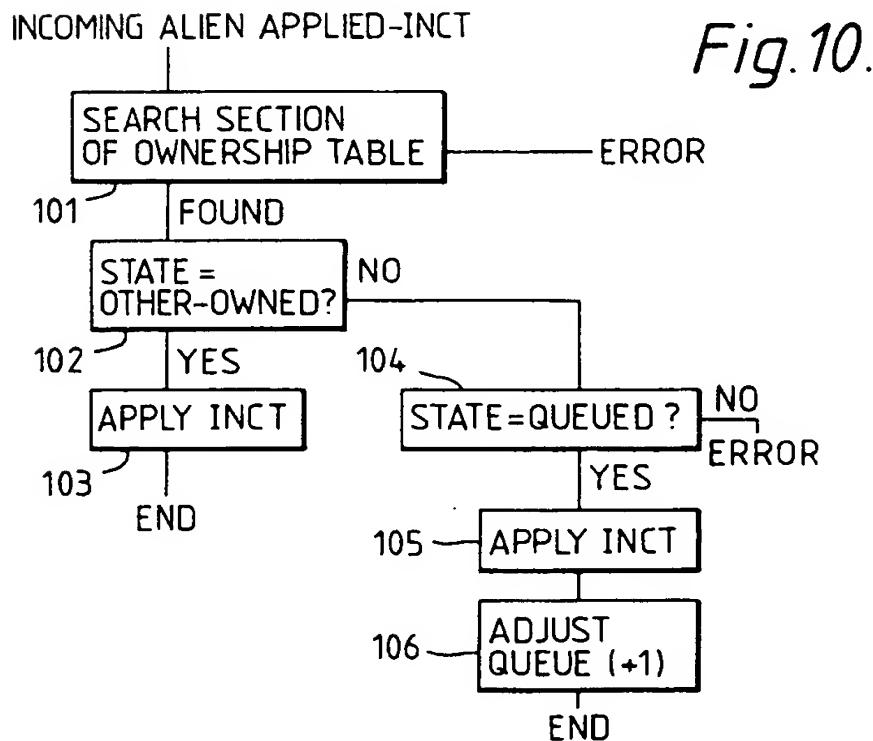
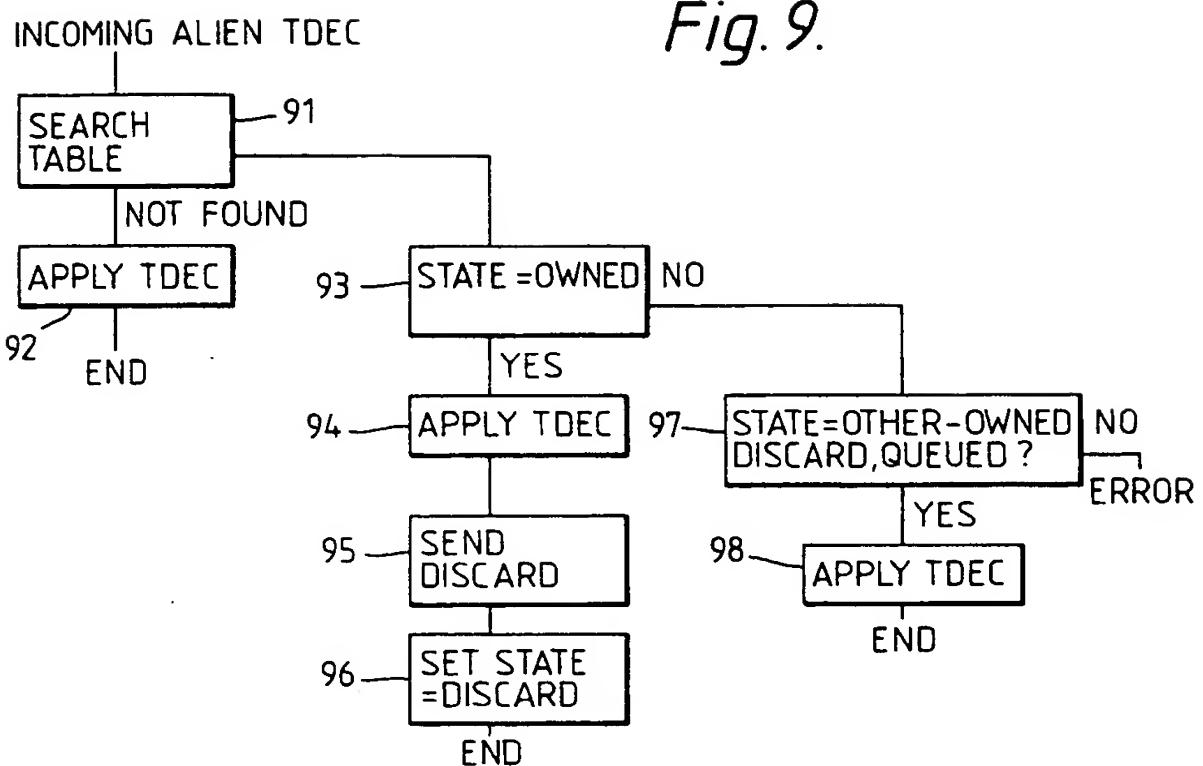
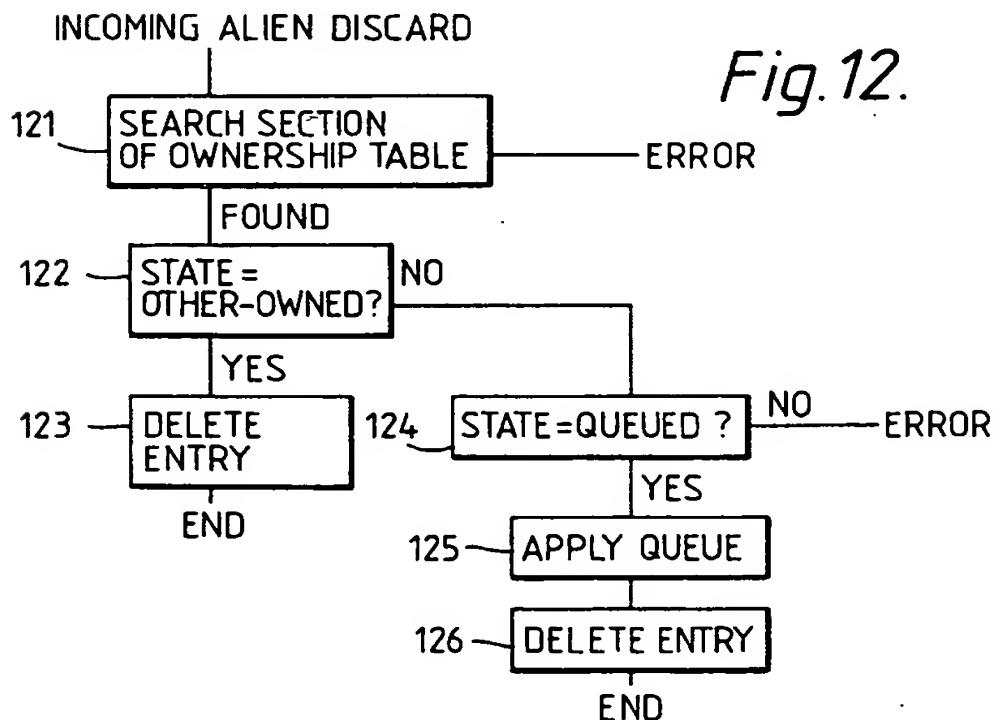
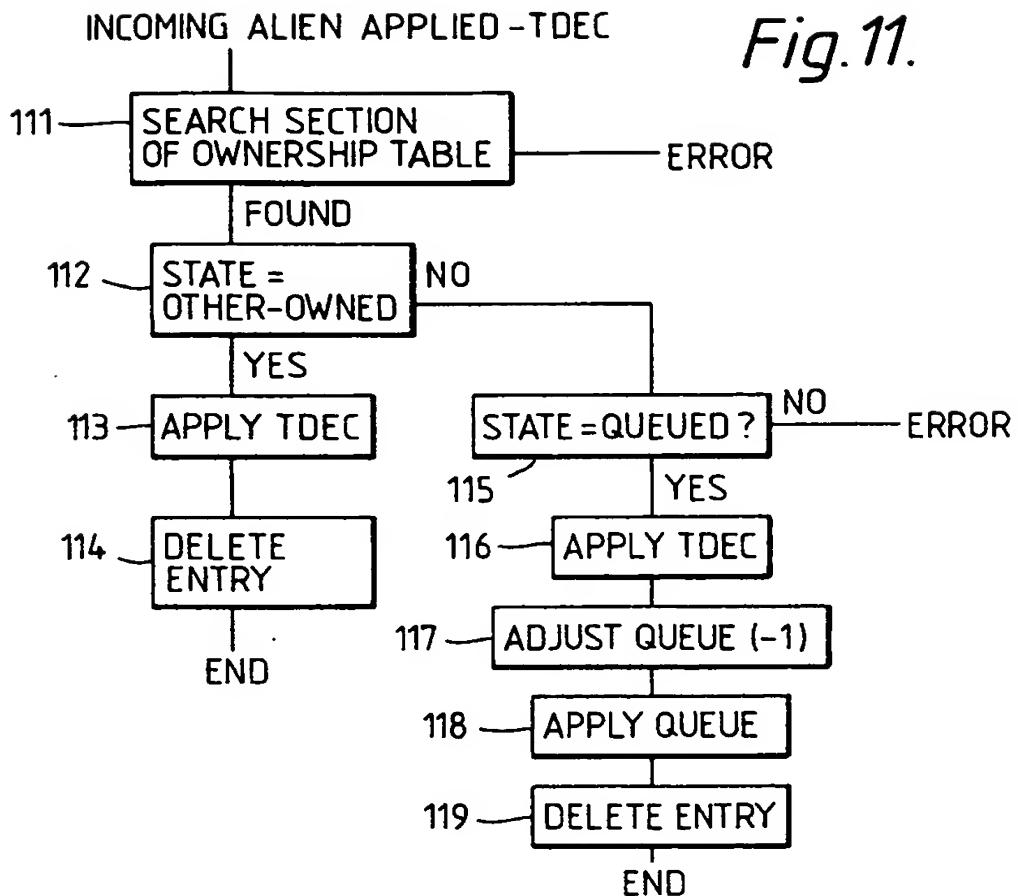


Fig. 8.









European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 92 30 8365

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. CL.S)
Y	EP-A-0 351 556 (MODULAR COMPUTER SYSTEMS INC.) * the whole document * ---	1-7	G06F9/46
Y	EDN - ELECTRICAL DESIGN NEWS vol. 36, no. 3, 4 February 1991, NEWTON, MA, US pages 85 - 90. D. L. RIPPS 'Semaphores and controlled shared variables' * page 86, left column, line 4 - page 88, left column, line 14 * ---	1-7	
A	EP-A-0 013 301 (INTERNATIONAL BUSINESS MACHINES CORPORATION) * the whole document * ---	1-7	
TECHNICAL FIELDS SEARCHED (Int. CL.S)			
G06F			
The present search report has been drawn up for all claims			
Place of search	Date of completion of the search	Examiner	
BERLIN	.08 FEBRUARY 1993	R. ABRAM	
CATEGORY OF CITED DOCUMENTS			
X : particularly relevant if taken alone	T : theory or principle underlying the invention		
Y : particularly relevant if combined with another document of the same category	E : earlier patent document, but published on, or after the filing date		
A : technological background	D : document cited in the application		
O : non-written disclosure	L : document cited for other reasons		
P : intermediate document	& : member of the same patent family, corresponding document		

This Page Blank (uspto)

This Page Blank (uspto)